

16/670,998 PTO-892

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
29 September 2005 (29.09.2005)

PCT

(10) International Publication Number
WO 2005/091211 A1

(51) International Patent Classification⁷: G06K 9/00

(21) International Application Number:
PCT/US2005/008735

(22) International Filing Date: 16 March 2005 (16.03.2005)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/554,050 16 March 2004 (16.03.2004) US
10/965,687 13 October 2004 (13.10.2004) US
00/000,000 15 March 2005 (15.03.2005) US

(71) Applicant (for all designated States except US): 3VR SECURITY, INC. [US/US]; 185 Berry Street, Suite 4150, San Francisco, California 94107 (US).

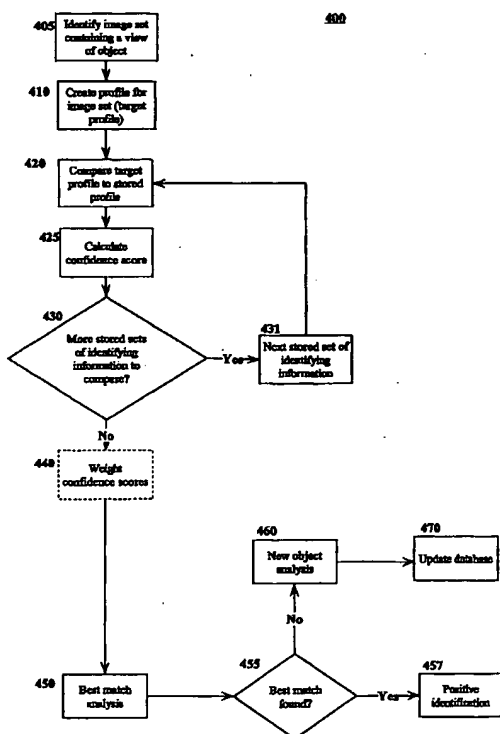
(72) Inventors; and
(75) Inventors/Applicants (for US only): HAUPT, Gordon, T., PH.D. [US/US]; 509 29th Street, San Francisco, California 94131 (US). FREEMAN, J. Andrew [US/US]; 75 South 13th Street, San Jose, California 95112 (US). FLEISCHER, Stephen, D., PH.D. [US/US]; 175 Bluxome Street, #120, San Francisco, California 94107 (US). VALLONE, Robert, P. [US/US]; 2626 Bryant Street, Palo Alto, California 94306 (US). RUSSELL, Stephen, G. [US/US]; 108 San Jose Street #3, San Francisco, California 94107 (US). FREDERICK, Timothy B. [US/US]; 226 Grand View Avenue, San Francisco, California 94114 (US).

(74) Agents: BOVERI, Lesley, C. et al.; Hickman Palermo Truong & Becker, 2055 Gateway Place, Suite 550, San Jose, CA 95110 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BR, BG, BR, BW, BY, BZ, CA, CH, CN,

[Continued on next page]

(54) Title: INTERACTIVE SYSTEM FOR RECOGNITION ANALYSIS OF MULTIPLE STREAMS OF VIDEO



(57) Abstract: A method of identifying an object captured in a video image in a multi-camera video surveillance system is disclosed. Sets of identifying information are stored in profiles, each profile being associated with one object. The disclosed method of identifying an object includes comparing identifying information extracted from images captured by the video surveillance system to one or more stored profiles. A confidence score is calculated for each comparison and used to determine a best match between the extracted set of identifying information and an object. In one embodiment, the method is used as part of a facial recognition system incorporated into a video surveillance system.

WO 2005/091211 A1



CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

FR, GB, GR, HU, IE, IS, IT, LT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

- with international search report
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI,

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

INTERACTIVE SYSTEM FOR RECOGNITION ANALYSIS OF MULTIPLE STREAMS OF VIDEO

FIELD OF THE INVENTION

[0001] The present invention relates to recognizing or identifying objects from images taken in naturalistic environments and, more specifically, to a system that improves accuracy in facial recognition by allowing a user to correct and update selections made by a facial recognition module such that multiple sets of identifying information can be associated with a single person and used by the facial recognition module to improve future matching.

BACKGROUND

[0002] "Biometrics" refers to unique physiological and/or behavioral characteristics of a person that can be measured or identified. Example characteristics include height, weight, fingerprints, retina patterns, skin and hair color, and voice patterns. Identification systems that use biometrics are becoming increasingly important security tools. Identification systems that recognize irises, voices or fingerprints have been developed and are in use. These systems provide highly reliable identification, but require special equipment to read the intended biometric (e.g., fingerprint pad, eye scanner, etc.) Because of the expense of providing special equipment for gathering these types of biometric data, facial recognition systems requiring only a simple video camera for capturing an image of a face have also been developed.

[0003] In terms of equipment costs and user-friendliness, facial recognition systems provide many advantages that other biometric identification systems cannot. For instance, face recognition does not require direct contact with a user and is achievable from relatively far distances, unlike most other types of biometric techniques, e.g., fingerprint and retina pattern. In addition, face recognition may be combined with other image identification methods that use the same input images. For example, height and weight estimation based on comparison to known reference objects within the visual field may use the same image as face recognition, thereby providing more identification data without any extra equipment.

[0004] However, facial recognition systems can have large error rates. In order to provide the most reliable and accurate results, current facial recognition systems typically

face recognition does not require direct contact with a user and is achievable from relatively far distances, unlike most other types of biometric techniques, e.g., fingerprint and retina pattern. In addition, face recognition may be combined with other image identification methods that use the same input images. For example, height and weight estimation based on comparison to known reference objects within the visual field may use the same image as face recognition, thereby providing more identification data without any extra equipment.

[0005] However, facial recognition systems can have large error rates. In order to provide the most reliable and accurate results, current facial recognition systems typically require a person who is to be identified to stand in a certain position with a consistent facial expression, facing a particular direction, in front of a known background and under optimal lighting conditions. Only by eliminating variations in the environment is it possible for facial recognition systems to reliably identify a person. Without these types of constraints in place, the accuracy rate of a facial recognition system is poor, and therefore facial recognition systems in use today are dedicated systems that are only used for recognition purposes under strictly controlled conditions.

[0006] Video surveillance is a common security technology that has been used for many years, and the equipment (i.e., video camera) used to set up a video surveillance system is inexpensive and widely available. A video surveillance system operates in a naturalistic environment, however, where conditions are always changing and variable. A surveillance system may use multiple cameras in a variety of locations, each camera fixed at a different angle, focusing on variable backgrounds and operating under different lighting conditions. Therefore, images from surveillance systems may have various side-view and/or top-view angles taken in many widely varying lighting conditions. Additionally, the expression of the human face varies constantly. Comparing facial images captured at an off-angle and in poor lighting with facial images taken at a direct angle in well lit conditions (i.e., typical images in a reference database) results in a high recognition error rate.

[0007] In a controlled environment, such as an entry vestibule with a dedicated facial recognition security camera, the comparison of a target face to a library of authorized faces is a relatively straightforward process. An image of each of the authorized individuals will have been collected using an appropriate pose in a well lighted area. The person requesting entry to the secured facility will be instructed to stand at a certain point relative to the

camera, to most closely match the environment in which the images of the authorized people were collected.

[0008] For video surveillance systems, however, requiring the target individual to pose is an unrealistic restriction. Most security systems are designed to be unobtrusive, so as not to impede the normal course of business or travel, and would quickly become unusable if each person traveling through an area were required to stop and pose. Furthermore, video surveillance systems frequently use multiple cameras to cover multiple areas and especially multiple entry points to a secure area. Thus, the target image may be obtained under various conditions, and will generally not correspond directly to the pose and orientation of the images in a library of images.

[0009] The approaches described in this section are approaches that could be pursued, but not necessarily approaches that have been previously conceived or pursued. Therefore, unless otherwise indicated, it should not be assumed that any of the approaches described in this section qualify as prior art merely by virtue of their inclusion in this section.

SUMMARY OF EMBODIMENTS OF THE INVENTION

[0010] Techniques are provided for improving accuracy of an object recognition system in a naturalistic environment. These techniques may be used, for example, for providing accurate facial recognition in a video surveillance system.

[0011] In one embodiment, a method is provided for determining a best match between a target profile and a set of stored profiles, where a profile contains a set of identifying information extracted from an image set associated with the profile. The method includes generating a plurality of confidence scores based on comparisons between the target profile and the set of stored profiles. The generated confidence scores are weighted using information external to the confidence scores. Based on the plurality of weighted confidence scores, a stored profile is selected as the best match for the target profile.

[0012] In one embodiment, a method for maintaining associations between profiles and objects in an object recognition system is provided. The method includes automatically creating an association between a first stored profile and a first object, and automatically creating an association between a second stored profile and the first object. Views of the image sets associated with the first and second stored profiles are provided to a user. Feedback is received from the user about the association between the second stored profile

and the first object. The second stored profile's association with the first object is modified in accordance with the received feedback.

[0013] In one embodiment, a method is provided for determining a best match between a target profile and an object in an object recognition system in which each object recognized by the system is associated with a plurality of stored profiles. A profile contains a set of identifying information extracted from an image set associated with the profile. The method includes generating a plurality of confidence scores based on comparisons between the target profile and a set of stored profiles. The object associated with each generated confidence score is determined. The generated confidence scores and determined associated objects are analyzed, and the analysis is used to select a best matching object in accordance with the analyzed confidence scores and determined associated objects.

BRIEF DESCRIPTION OF THE DRAWINGS

[0014] The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

[0015] FIG. 1 is a block diagram that illustrates a multi-camera video processing pipeline architecture upon which an embodiment of the invention may be implemented;

[0016] FIG. 2 is a block diagram that illustrates a computer system upon which an embodiment of the invention may be implemented;

[0017] FIG. 3 is a diagram that illustrates a high-level view of a video surveillance network on which an embodiment of the invention may be implemented;

[0018] FIG. 4 is a flowchart illustrating one embodiment of a method for processing video data in a multi-camera image recognition system;

[0019] FIG. 5 is a flowchart illustrating one embodiment of a method for performing New Object Analysis; and

[0020] FIG. 6 is a diagram illustrating one example of the relationship between feature sets and known persons.

DETAILED DESCRIPTION OF EMBODIMENTS OF THE INVENTION

[0021] In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, that the present invention may be practiced without these specific

details. In other instances, well-known structures and devices are shown in block diagram form in order to avoid unnecessarily obscuring the present invention.

[0022] A surveillance system that includes face recognition capabilities to identify people in the video images acquired by the surveillance system could have many important applications. For example, such a system could be used as a security system to grant or deny access to select individuals, to sound an alarm when a particular person is recognized, or to continuously track an individual as the individual travels amongst a plurality of people, and so forth.

[0023] In order to incorporate facial recognition into a video surveillance system, however, it must be possible to identify a single frame or video clip that may contain an image of a person's face, extract identifying information from the image for comparison with known faces, and reliably determine whether the extracted identifying information matches identifying information of the face of a known person. Current surveillance systems with face recognition capabilities have not successfully been able to perform each of these steps.

[0024] As discussed above, high accuracy rates for facial recognition have only been achieved in systems operating under very controlled conditions, as the accuracy of matches found by a facial recognition system dramatically decreases with changes in a subject's face orientation, changes in illumination conditions, and changes in a subject's facial expressions. These limitations mean that the use of facial recognition has been limited to access control points where a cooperative subject is standing still, facing the camera, and lighting is controlled. These fundamental restrictions prevent current face recognition systems from effectively and reliably identifying individuals in field-deployable conditions. Only if both the surveillance image and the reference image are taken from the same angle, and with consistent lighting and facial expression, is any type of significant accuracy achieved. Video surveillance systems that operate in a naturalistic environment in which subjects are not required to pose for an identification camera under controlled circumstances may have an accuracy rate so low as to make the system unusable.

[0025] Furthermore, in an image tracking system, wherein an image of a target is obtained from one scene, and then matched to images in subsequent scenes, neither the original image nor the subsequent images will be obtained under ideal conditions, thereby reducing recognition accuracy rates even further.

[0026] Over time, people will routinely change hairstyles, hair color, suntan, makeup, posture, and body weight. Moreover, the facial characteristics of a person's face will change due to aging. These types of changes make it difficult for a facial recognition system to consistently identify a person correctly using an unchanging static image of a person as its reference, even if other environment variables are controlled.

[0027] A system which provides for reliable facial recognition in a multi-camera naturalistic environment, such as a surveillance system, is disclosed herein. In addition, the disclosed techniques can be used to "train" the surveillance system to more accurately identify people over time.

[0028] Embodiments of the present invention provide for using images from a multi-camera surveillance system to construct identifying information about the objects or persons regularly surveyed by the system. Significantly, the cameras in the surveillance system may be operated under different and variable lighting conditions, and with various zoom and focus settings. In a facial recognition system embodiment, the faces of people captured by the cameras in the system can be facing different directions, have various facial expressions, and will change over time. Further embodiments of the present invention may be used to recognize any type of object from images acquired by a multi-camera, naturalistic environment surveillance system.

[0029] Unlike other facial recognition systems, a surveillance system has few artificial constraints, and provides a large number of naturalistic images of people or objects over time. Embodiments of the present invention use the images from a multi-camera surveillance system to acquire a large number of naturalistic positive identifications of an object or person, and use these positive identifications to train the system to improve recognition rates.

EXEMPLARY SYSTEM

[0030] FIG. 3 illustrates a high-level pictorial representation of a video surveillance network 300 in which an embodiment of the present invention may be implemented. As shown, video cameras 310, 312, 314 are connected to network 302, as are voice recorder 318, server 340, expert user workstation 330 and storage unit 360. Network 302 is the medium used to provide communications links between various devices and computers connected together within the system. Surveillance network 302 may be implemented as any type of network, such as an intranet, a local area network (LAN), or a wide area network (WAN).

The network 302 may also comprise secure connections to the Internet. Network 302 may include connections such as wire, wireless communication links, or fiber optic cables, for example. Alternatively, instead of a network, some or all of the components of the surveillance system may be directly connected to each other.

[0031] Extraction module 352 extracts identifying information from the video data produced by cameras 310, 312, 314, and/or from samples taken by voice recorder 318. Extraction module 352 may use any method known to those skilled in the art that takes raw image data and extracts identifying information. Extraction module can be a component provided by a third party that integrates with system 300.

[0032] Matching module 350 processes the extracted identifying information produced by extraction module 352 to determine if the extracted identifying information matches identifying information stored in storage unit 360. Matching module 350 may use any method known to those skilled in the art to compare extracted sets of identifying information, such as feature sets for face recognition, to stored sets of identifying information, and calculate a "confidence score." A confidence score numerically represents the similarity between target identifying information extracted by module 352 and a stored set of identifying information. Matching module 350 may also be a component provided by a third party that integrates with system 300.

[0033] As shown in FIG. 3, extraction module 352 and matching module 350 may be components of server 340. Alternatively, one or both of these modules may reside in a separate computer dedicated to performing just that module.

[0034] Video cameras 310, 312 and 314 may be any cameras used in a video surveillance system, either visible or hidden from persons surveyed by the surveillance system. Cameras 310, 312, 314 may operate in the visual range of the electromagnetic spectrum or may include other ranges including infrared (IR) and ultraviolet (UV). In addition, a camera may also have light amplification capabilities for low light conditions. Cameras 310, 312, 314 may be identical, or each may have different capabilities.

[0035] Voice recorder 318 may be used in conjunction with the images acquired by cameras 310, 312, 314 to identify a person. While shown in the example embodiment of FIG. 3, voice recorder 318 is not required. Likewise, while only one voice recorder 318 is shown in FIG. 3, any number of voice recorders could be used.

[0036] Data store 360 may contain one or more databases of video data recorded by cameras 310, 312, 314. Video data stored in data store 360 may include single frames or images, as well as video clips. Data store 360 may also include one or more databases of audio or voice samples captured by the surveillance system. In addition, data store 360 may also contain one or more reference databases of identifying information associated with objects or persons whose image was obtained by a camera in the multi-camera surveillance system. Data store 360 may contain additional databases that store surveillance system management information. Data store 360 may be one device. Alternatively, each database in data store 360 may be stored on separate storage devices in separate locations. Data store 360 is intended to include any computer that stores video data and surveillance system management information. Video data stored in the system may include video data captured by cameras in the surveillance system, or may originate outside of the surveillance system. Data store 360 is accessible by matching module 350 to compare images acquired by any camera 310, 312, 314, and identifying information extracted from these images, to identification information and images stored in a database on data store 360.

[0037] Surveillance system 300 may include additional detection means, servers, clients and other peripheral devices not shown. For example, surveillance system 300 may also include Radio Frequency identification transponders used to identify individuals or objects to which the transponder is attached. FIG. 3 is intended as an example, and not as an architectural limitation for the present invention.

PIPELINE ARCHITECTURE

[0038] One specific example of multi-camera architecture that could be used to implement an embodiment of the present invention is disclosed in co-pending application U.S. Patent Application Serial No. 10/965,687, entitled PIPELINE ARCHITECTURE FOR ANALYZING MULTIPLE STREAMS OF VIDEO, filed on October 13, 2004, the contents of which have been incorporated by reference in their entirety for all purposes. FIG. 1, taken from the referenced co-pending application, herein after referred to as the "Pipeline Application", illustrates an embodiment of the multi-camera pipeline architecture.

[0039] In the system disclosed in the co-pending Pipeline application, numerous video analysis applications can access and analyze video data that represents video streams flowing through the pipeline, and annotate portions of the video data (e.g., frames and groups of

frames), based on the analyses performed, with information that describes the portion of the video data. These annotations flow through the pipeline, possibly along with corresponding frames or groups of frames, to subsequent stages of processing, at which increasingly complex analyses can be performed. Analyses performed at the various stages of the pipeline can take advantage of the analyses performed at prior stages of the pipeline through use of the information embodied in the annotations. At each stage of the pipeline, portions of the video streams determined to be of no interest to subsequent stages are removed from the video data, which reduces the processing requirements of the subsequent stages.

[0040] Ultimately, "events" are constructed and stored in a database, from which cross-event and historical analyses may be performed and associations with, and among, events may be made. Such events contain whatever information is relevant to describing the real-world activities or objects for which the event was constructed to describe. In addition, events may contain pointers to locations in persistent memory, e.g., a file store in storage unit 360 of FIG. 3, at which the associated frames and/or groups of frames are stored. Hence, from an event stored in the database, the associated frames and/or groups of frames can be replayed for further human-based or application-based analyses.

[0041] In one embodiment, the pipeline comprises four different successive stages of processing: (1) quick frame processing; (2) deep frame processing; (3) cluster processing; and (4) database processing. Due to the nature of the pipeline, applications plugged into the pipeline, via application program interfaces (APIs) associated with each respective stage, can perform increasingly more complex analyses at each successive stage of processing.

[0042] Generally, as the videos flow down the pipeline, (1) portions of the videos or frames that are considered uninteresting to all the applications at a given stage are removed, thereby reducing the amount of data that flows further down the pipeline; (2) portions of the videos or frames that are considered interesting to an application at a given stage are analyzed, with a goal of identifying features, activities, objects, etc. of interest; and (3) analyzed portions of the videos or frames are annotated by the applications with information that describes what the applications identified as interesting in that portion of the video.

[0043] Stage 1 of the pipeline processing ("P1") is referred to as "quick frame" processing. Quick frame processing is lightweight processing (i.e., not relatively resource-intensive and computationally complex) performed in real-time as the video streams flow

into the pipeline. Various applications tailored to recognize and identify certain occurrences may plug into the pipeline via the quick frame API, to perform fast lightweight operations, such as noise reduction, motion detection, gross object finding (e.g., a vehicle or person), object tracking, frame area filtering, and the like.

[0044] Stage 2 of the pipeline processing ("P2") is referred to as "deep frame" processing. Any number of video analysis applications (referred to in FIG. 1 as P2 Analyzers, P2A1, P2A2, P2An) can access the video feeds from the buffers for deep frame processing, through a deep frame API. Various applications tailored to recognize and identify certain occurrences may plug into the pipeline via the deep frame API, to perform more computationally complex and resource-intensive analyses operations than with quick frame processing. For example, deep frame analyses of the video streams may include face finding, license plate recognition, complex object detection (e.g., gun finding), feature extraction, and the like.

[0045] An application to identify a frame from a camera that contains an image of a person's face may be included as a "P2" application that uses the deep frame API. Likewise, an application to extract features from faces identified in frames of surveillance video from multiple cameras (e.g., extraction module 352 of FIG. 3) may also be included as a "P2" application. When the pipeline architecture is used for face recognition, a preliminary identification of a person in a single video frame or image may also be made by an application using the P2 API.

[0046] If any P2 analyzer finds particular frames to be of interest, then the analyzer determines what type of analysis to perform on the video clip or frame, and creates "pipeline objects" based thereon. A pipeline object herein refers to a programmatic object such as an object in object-oriented programming. Pipeline objects created at the deep frame processing stage typically contain a pointer to one or more relevant frames, and additional information about the content of the frame on which the pipeline object is based. Sequences of pipeline objects are output from the deep frame processing stage and, in one embodiment, are queued in a buffer between the deep frame processing stage and the cluster processing stage.

[0047] Stage 3 of the pipeline processing ("P3") is referred to as "cluster" processing. Any number of video analysis applications (referred to in FIG. 1 as P3 Analyzers, P3A1, P3A2, P3An) can access the video feeds and other information from buffers for cluster and

event processing, through a cluster API. Various applications tailored to recognize and identify certain occurrences may plug into the pipeline via the cluster API, to perform analyses on the video streams across time (i.e., across frames) and across cameras (i.e., within a "cluster" of cameras that, for analysis purposes, are treated as an entity). Events based on analyses of the video streams at the cluster stage of processing may include various tailored analyses and construction of associated events, such as person or face events, alert generation events, externally triggered events, and the like.

[0048] An event that is constructed based on video feeds from multiple cameras, i.e., a cluster of cameras, is referred to as a "cluster event." Cluster events provide information such as what happened in a building lobby rather than what happened in view of camera X, where camera X is only one of a plurality of cameras operating in the lobby.

[0049] The same video data can be used in multiple stages of processing. For example, a P2 application can be used to make a preliminary identification of a face captured in a single image or frame. Then, during P3 processing a cluster event may be created that includes the frame used to make the preliminary identification. A P3 application can be used to link multiple related images from a cluster event into an image set, and then analyze the image set to identify the person, resulting in a more reliable identification than the preliminary identification made by a P2 application.

[0050] Events, either cluster events or non-cluster events, are constructed by P3 analyzers at the cluster stage of processing, based on output by the deep frame stage of processing. Events are output from the cluster stage and stored in a database. In one embodiment, each event is embodied as a row in a database table, where each row contains (1) information that describes whatever the analyzer determined about what occurred in the area observed (i.e., the content of the video frames or video clips), for which the event was constructed, and (2) references to the frames or video clips that are associated with the event, if desired or necessary, including pointers to the frames or video clips in a file store. The P3 analyzer applications determine what information to store in the database in association with an event.

[0051] Further analysis and reasoning can be applied to events, or combinations of events, that are stored in the database. From a database record containing pointers to the location in the file store at which frames and video clips are stored, the associated frames and

video clips can be replayed and reviewed, for example, by a user via a display monitor or by database stage analyzer applications via a database API.

[0052] Stage 4 of the pipeline processing ("P4") is referred to as database processing. Any number of video analysis applications (referred to in FIG. 1 as P4 Analyzers, P4A1, P4A2, P4An) can access event records from the database for database processing, through the database API. Various applications tailored to perform complex analysis across events and across clusters may plug into the pipeline via the database API, to perform analyses such as historical analyses, person/place/time reports, object identification, and the like. As discussed in more detail below, the New Object Analysis in which an expert user makes a positive identification of a person or object in a video image may be a P4 application that uses the database API.

[0053] The above-described Pipeline architecture may be used to implement embodiments of the techniques described hereafter, although as will be apparent to those skilled in the art, embodiments may be implemented in any multi-camera surveillance system, and are not limited to this architecture.

EXTRACTING A SET OF IDENTIFYING INFORMATION

[0054] The techniques disclosed herein are described using facial recognition as an example application, however, the techniques are not limited to just facial recognition. The disclosed techniques may be used to recognize and identify any object whose image is obtained in a multi-camera surveillance system, such as a weapon, suitcase, vehicle and the like. Furthermore, although the techniques are described using video cameras, it will be apparent to those skilled in the art that any camera or device used to produce a sample, such as an image, can be used. For example, voice samples may be recorded from multiple recorders and used as identification input.

[0055] Most facial recognition systems do not directly compare images to effect a recognition. Instead, each face is characterized using a predefined set of characteristic parameters, such as the ellipticity of the face, the spacing of the eyes, the shape of the chin, etc. A search for a match to a reference face is based on a comparison of these characteristic parameters instead of directly comparing images. These characteristic parameters are designed to facilitate a distinction between images of different faces, and a matching between

different images of the same face. In this manner, the characteristic parameters of a target image can be compared to the characteristic parameters of a reference image.

[0056] Typically, in facial recognition systems, the set of characteristic parameters is called a "feature set." A feature set for a person's face captured in an image may contain mathematical expressions or vectors that represent various facial profile measurements or correspond to certain facial features. As is known to those skilled in the art, there are many different known types of facial feature sets that can be created, and the present invention is not limited to any one type of facial feature set. In addition, while embodiments of the present invention are described using a facial recognition system as an example, alternative embodiments of the present invention may identify non-person objects by using characteristic extraction parameters related to the type of object being identified.

[0057] In one embodiment, a "profile" contains a set of identifying information associated with a view of the object shown in an image or set of images. For example, a profile of a person may contain a feature set extracted from a view of a person's face in an image or set of images.

[0058] If a profile is created from a single image, the set of identifying information in the profile is extracted from that single image. If a profile is created from a set of multiple images, the set of identifying information in the profile may be calculated a number of ways. For example, the profile's set of identifying information may contain identifying information extracted from the image in the set of images that is determined to contain the "best" view of the object. As another example, the profile's set of identifying information may be calculated by averaging sets of identifying information extracted from each image in the image set. As another example, the profile's set of identifying information may be extracted from an image that is created by averaging the images in the image set. The profile's set of identifying information may include multiple subsets of identifying information, each subset of identifying information extracted from an individual image. Any method of extracting and calculating a set of identifying information from a set of images may be used to create the set of identifying information that is stored in a profile associated with that image set.

[0059] A profile may optionally contain other information in addition to the set of identifying information, such as identification of the camera(s) used to capture the associated image(s), or time and/or location information, for example. By including additional

information in a profile, multiple profiles can be associated with the same person or object as each profile represents a separate occurrence of that person or object captured by the surveillance system.

MULTIPLE PROFILES CAN BE ASSOCIATED WITH AN OBJECT

[0060] Significantly, embodiments of the present invention allow multiple profiles to be associated with one person or object, where different profiles for the same object may include sets of identifying information that have different measurements for the same characteristic. For example, a first profile may contain a feature set for a particular person based on a first view of the person, and therefore reflect different characteristics than the feature set in a second profile for the same person taken under different conditions.

[0061] FIG. 6 illustrates this concept. In the example shown in FIG. 6, each person has a one-to-many relationship with stored profiles. Specifically, each person may have many profiles, but each profile can only be associated with one person.

[0062] In the example system of FIG. 6, a profile database in storage unit 360 contains ten profile records A1, A2, A3, B1, B2, B3, C1, C2, C3 and C4. Three people (Adam, Bart and Charlie) have previously been positively identified. These prior identifications are reflected in a database that associates a profile with a person. As shown, three profiles are associated with Adam (A1-A3), three profiles are associated with Bart (B1-B3) and four profiles are associated with Charlie (C1-C4).

[0063] Typically, each profile associated with a person reflects measurements produced in a different manner than the measurements of other profiles for the same person. For example, the different profiles for a particular person may have been derived from different sources. Thus, profiles A1, A2 and A3 associated with Adam may have been derived from three different images of Adam. Alternatively, profiles A1, A2 and A3 may have been derived from the same image of Adam, but may have used different algorithms to derive measurements of the same characteristics.

[0064] When a facial recognition system using the techniques described herein performs a matching or identification analysis, it is possible that multiple profiles for the same person will be determined to be potential matches. The occurrence of multiple profiles as likely matches in a matching analysis may be an indicator that the person associated with the multiple matching profiles is a "best match", as will be discussed in more detail below.

[0065] Known recognition systems are typically constrained to identifying a single person or object as a match for an object in a target image. Typically, such systems constrained to this "Yes/No" analysis calculate confidence scores between a target image and reference images. The person associated with the highest-ranked confidence score is identified as the person in the target image, so long as the confidence score for the match is greater than a minimum confidence threshold level. Otherwise, the system will indicate that no matches were found. The accuracy rates in these systems may be especially low. These systems will have a high rate of both misidentification and non-identification. A misidentification occurs when the wrong person is identified, and is sometimes called a "false positive", or "false acceptance." A non-identification occurs when a match is not found, although the person is in the system, and is sometimes called a "false negative" or "false rejection."

[0066] Embodiments of the present invention perform further analysis on the confidence scores, and determine a "best match", rather than simply determining a highest-ranked confidence score, resulting in significantly higher recognition rates.

[0067] In one embodiment, the confidence scores calculated by a matching module may be weighted using external information, and the highest weighted score chosen as the best match. In another embodiment, as objects can be associated with multiple profiles, a weighted average of confidence scores may be calculated for each object, and the object with the highest weighted average chosen as the best match. In another embodiment, weighted averages of weighted confidence scores may be used to determine the best match. Generally, the techniques described herein can use many types of information external to the confidence score generated by a matching module to weight the confidence scores or perform weighted averaging of confidence scores such that a more reliable identification can be made. For example, embodiments can use such information as physical location of an object within an image, time proximity of an object in an image to another object that has been positively identified, or organizational or group information related to potential matching candidate objects, to weight the confidence scores or perform weighted averaging. Any type of external information can be used, and the disclosed techniques are not limited to the examples given herein.

OBJECT RECOGNITION PROCESS

[0068] Referring to FIG. 4, a flowchart illustrating an object recognition process using weighted confidence scores or weighted averaging is depicted in accordance with one embodiment of the present invention. In one embodiment, process 400 is started when an event trigger occurs. For example, if implemented using the architecture disclosed in the co-pending Pipeline Application, one "trigger" that may start process 400 could be detection of a face in surveillance video by a "P2" face finding application that uses the P2 pipeline API.

[0069] The process begins at step 405 when an image set that contains a view of the object of interest is identified. The image set may include just a single image from a single camera. Alternatively, the image set may include multiple images from a single camera, or multiple images from multiple cameras. For example, an image set may be identified as a set of images related to a cluster event. The term "image set" will be used herein to include a set that comprises a single image as well as a set of multiple images.

[0070] As discussed, one common object recognition system is a facial recognition system. However, embodiments of the present invention may be used to recognize any type of object in an image set, and the image set identified at step 405 may contain a view of any type of object. When the techniques disclosed herein are used in a facial recognition system, the identified image set contains image(s) of a person's face as captured in frame(s) of video surveillance camera(s).

[0071] At step 410, a profile for the object in the image set is created. In this step, identifying information is extracted from the image set. In an embodiment of the present invention that uses the pipeline architecture of the Pipeline Application, third-party software can use the P2 and/or P3 APIs to extract a set of identifying information for an image set, and the created profile includes this extracted set of identifying information. Additional information, such as camera identification(s), date, time, etc., may also be included in the created profile.

[0072] The profile created in step 410, referred to herein as the "target profile", is compared against stored profiles. At step 420, the set of identifying information in the target profile ("target identifying information") is compared to a set of identifying information from a profile stored in a reference database. A confidence score that numerically represents the similarity between the target identifying information and the set of identifying information in the stored profile is calculated in step 425. In an embodiment of the present invention that

uses the pipeline architecture of the Pipeline Application, third-party software can use the P2 and/or P3 APIs to make the comparison and generate a confidence score.

[0073] In a facial recognition system embodiment, the profile is compared to profiles stored in a database containing a plurality of stored profiles, and/or images. Typically, the stored profiles are associated with a known, identified person or object, but this is not required. For example, the comparison may compare the target profile to a stored profile that contains a feature set taken from earlier video capturing the face of an intruder who has not yet been identified.

[0074] Loop 430-431-420-425 is performed for each stored profile that is to be compared to the target profile. In particular, a confidence score is created for each comparison. In a preferred embodiment, a comparison is made with each stored profile, although it is possible that comparisons will only be made with a selected subset of the stored profiles.

[0075] While known facial recognition systems will simply return the top-ranked match when at least one confidence score exceeds a minimum threshold, the techniques described herein perform further analysis on the confidence scores to determine the best match.

WEIGHTED CONFIDENCE SCORES

[0076] After all comparisons have been made, in the embodiment shown in FIG. 4, the confidence scores generated by loop 420-425-430-431 may be weighted at optional step 440. Alternatively, the confidence scores may be weighted in the loop as comparison scores are generated. Step 440 is optional, as confidence scores do not have to be individually weighted.

[0077] There are many methods of weighting confidence scores contemplated, and several non-limiting examples are given herein. Generally, the weighting factor is derived from information external to the confidence scores themselves. For purposes of discussion, consider an example scenario in which a target profile is compared to five stored profiles, and five confidence scores are calculated, as shown in Table 1:

Stored Profile	Confidence Score
A	90
B	85
C	80

D	75
E	70

Table 1

[0078] If the confidence scores are used without weighting, profile A may be determined to be the closest match, and if the confidence score of 90 is greater than an object identification minimum threshold, the object in target profile may be identified as the object associated with profile A. As discussed above, error rates using this method can be quite high.

[0079] One example of weighting confidence scores such that a better match can be made is to weight the most recently created stored profiles higher, as characteristics of a person's face change over time, and more recent profiles may more closely reflect the person in the image. For example, profiles that are less than a week old may have a weighting factor of 1.0, profiles that are older than a week, but less than a month, may have a weighting factor of 0.9, and all other profiles are weighted by a factor 0.8. Assume, using the example from Table 1, that profile A is 3 weeks old, profile B is 1 hour old, profile C is 3 days old, profile D is 2 weeks old and profile E is 3 months old. Weighted confidence scores for this example are shown in Table 2:

Stored Profile	Confidence Score	Weight	Weighted Confidence Score
A	90	0.9	81
B	85	1.0	85
C	80	1.0	80
D	75	0.9	67.5
E	70	0.8	56

Table 2

[0080] In this example, even though profile A had the highest confidence score, profile B has the highest weighted confidence score, and may be selected as the closest matching profile for the target profile.

[0081] As another example, time proximity of the target image set to other image sets that contain positively identified objects may be used to weight the scores. That is, if a

person is positively identified in video caught by Camera 2 at time 04:25:03, then it is more likely that the person identified in an image captured by Camera 2 at time 04:26:30 is the same person. Time proximity across cameras may also be used as a weighting factor. For example, if Camera 1 in the lobby recognizes a person at time 02:30:00, then it is likely that the person captured by Camera 3 in the hallway connected to the lobby at time 02:32:00 is the same person.

BEST MATCH ANALYSIS USING WEIGHTED AVERAGING

[0082] Returning to FIG. 4, after the confidence scores are weighted, if they are weighted, a Best Match analysis occurs at step 450. Significantly, instead of simply returning the single highest ranked match, as in other object recognition systems, embodiments of the present invention may analyze the individual confidence scores (weighted or unweighted), and/or the persons or objects associated with the profiles that were used to calculate the confidence scores, to intelligently determine a best match in Best Match Analysis step 450.

[0083] There are many methods contemplated for performing a Best Match Analysis, and several non-limiting examples are given herein. In particular, it is contemplated that an average or weighted average of confidence scores can be used to determine a best match.

[0084] As an object may have multiple stored profiles associated with it, the number of profiles in a short list of the highest ranked profiles associated with each object may be considered. Because objects can be associated with multiple profiles, the short list of likely matches will include the correct person or object as a likely candidate in the list much more frequently than if the system is constrained to just selecting a single highest-ranked profile.

[0085] One simple example of a best match analysis technique that illustrates this concept is to select the person or object who has the greatest number of profiles with a confidence score (weighted or unweighted) that is greater than a "best match" minimum confidence level. FIG. 6 illustrates this concept. As shown in FIG. 6, person 601 is spotted in front of camera 310 in a multi-camera surveillance system. The face of person 601 is identified in a video image by a face finding application, and profile 601A is created for person 601 that includes a feature set extracted from the image by extraction module 352. Comparisons are made by matching module 350 between the target profile (601A) and profiles in a reference database 365. In this example, ten profiles A1-A3, B1-B3 and C1-C4

are compared to target profile 601A, and a confidence score is calculated for each comparison by matching module 350.

[0086] Suppose confidence scores are calculated by matching module 350 as shown in Table 3:

<u>Profile</u>	<u>Confidence Score</u>
A1	40
A2	45
A3	10
B1	95
B2	30
B3	50
C1	94
C2	93
C3	89
C4	91

Table 3

[0087] In this example system, a best match minimum confidence level established by the system is 90. Therefore the profiles whose confidence scores that meet this minimum are B1:95, C1:94, C2:93, C4:91. In this example, the list of candidate objects includes Bart and Charlie. Even though there are four profiles that exceed the best match minimum threshold, only two objects (Bart and Charlie) are associated with the profiles in the list of likely matching profiles.

[0088] The profile with the highest confidence score is B1, associated with Bart. However, in this example, even though the confidence score of the match with Bart is higher than any match to Charlie, because there are three matches to Charlie that are greater than the best match minimum confidence level whereas there is only one match to Bart, the system may determine that Charlie is the best match. In an alternative embodiment, no minimum threshold is needed and a best match is determined using all profiles. That is, any object associated with a stored profile used to calculate a confidence score is a candidate object.

[0089] Alternatively, an average or weighted average may be determined. For example, for every object associated with a stored profile that has a matching confidence score over a

certain threshold, all confidence scores for profiles associated with that object can be averaged. The object with the highest average confidence score may be determined the best match. In another alternative embodiment, there may be multiple matching modules that use different algorithms to return results, and all these results could be averaged together. Any type of weighted averaging may be used to determine a best match.

[0090] Alternatively, both scores and rankings can be used to calculate a weighted average for each candidate object, and the weighted average may be used to determine a best match. For purposes of discussion, consider an example scenario in which a target profile is compared to five stored profiles, five confidence scores are calculated, and each of the five stored profiles has been associated with (i.e., identified as) a person, as shown in Table 3:

Stored Profile	Confidence Score	Person
A	95	Adam
B	85	Bill
C	80	Bill
D	75	Adam
E	70	Bill

Table 3

[0091] There are many ways to calculate a weighted average for each candidate object (person) in Table 3 using confidence scores and/or rankings. Table 4 illustrates an example in which weights are assigned according to rank:

Stored Profile	Confidence Score	Person	Weight
A	95	Adam	1
B	85	Bill	0.75
C	80	Bill	0.5
D	75	Adam	0.25
E	70	Bill	0.125

Table 4

[0092] One technique is to add the weights for each candidate object, and not use the actual confidence scores. Using this technique, candidate object Adam would have a weighted average score of $1 + 0.25 = 1.25$. Bill would have a weighted average score of 0.75

+ 0.5 + 0.125 = 1.375. Using this example technique, Bill may be determined to be the best match.

[0093] Alternatively, the ranking weights could be used as multiplying factors. Using this technique, Adam would have a weighted average score of $[(95*1) + (75*0.25)]$, or 113.75. Bill would have a weighted average score of $[(85*0.75) + (80*0.5) + (70*0.125)]$, or 112.5. Using this technique, Adam may be determined to be the best match instead of Bill.

USING EXTERNAL INFORMATION IN BEST MATCH ANALYSIS

[0094] In addition, Best Match Analysis 450 may perform analysis using information external to the confidence scores and rankings. For example, the time that a candidate object was last identified by the surveillance system may be considered in the analysis. The external information is used to weight the averages of candidate objects.

[0095] As another example, organizational or other types of information associated with candidate objects may be factored into the analysis. For example, consider a facial recognition system installed at a secure office building that is connected to a system that stores information about the organization that occupies the building. In this example, work shifts for each potential candidate for the person in the image set may be looked up. A person who is scheduled to be in the building at the time the image set was captured may be considered a better match than a person who is not scheduled to work that shift, or who is scheduled to be on vacation. As another example, the authorized and/or likely locations for a person may be looked up. For example, whether a candidate works on the loading dock can be used when determining whether a target image taken at the loading dock is more likely to be that candidate. As another example, candidates who are identified as executives may be better matches for a person captured by an executive suite camera. The proximity of the person in the image set to other members of the same department may also be considered in the best match analysis.

[0096] As another example, a facial recognition system implemented using the techniques of the present invention may be connected to an access control system. In this case, the identification of a person as reported by an access control system can be used to weight the averages and determine a best match.

[0097] Any method of incorporating external information into the weights used to determine a best match may be used. In addition, various combinations of external information may also be used to weight the average, as well as combinations of external information, ranking and weighted confidence scores, as discussed above.

NEW OBJECT ANALYSIS

[0098] Returning to the embodiment shown in FIG. 4, if a best match cannot be determined, then New Object Analysis is performed on the image set in step 460. Using the example above, a configurable best match minimum confidence level may be set. If none of the weighted confidence scores is greater than this minimum confidence level, then it may be determined that no matches were found at step 455, and the process continues to the New Object Analysis step of 460.

[0099] Otherwise, the object identified by the Best Match Analysis is automatically identified as the object in the image by the system at step 457. In one embodiment, when an identification is made by the system in step 457, the profile created in step 410 is discarded. In alternative embodiments, the target profiles, or selected ones of the target profiles, are saved, at least for a period of time if not indefinitely. In one embodiment, target profiles can also be re-generated at a later time in order to perform Best Match Analysis or New Object Analysis again when more stored profiles are available for comparison, and thus result in more accurate identifications.

[0100] Although the embodiment shown in FIG. 4 shows that the New Object Analysis only occurs if no best match is determined, in alternative embodiments, New Object Analysis can occur at any time. For example, a system may be implemented such that New Object Analysis is performed for every object for a certain time period.

[0101] In New Object Analysis 460, an expert user can enroll a new user or correct/confirm a best match determination made by Best Match Analysis 450. Although FIG. 4 shows that expert review is only taken if no best match is found, in alternative embodiments, an expert user may review all identifications made by the system. That is, New Object Analysis may be used by an Expert User to override an incorrect identification made by the system.

[0102] New Object Analysis 460 can be performed at the time it is determined that an Expert User is needed to make the match. Alternatively, New Object Analysis 460 can be

performed at a later time, wherein the images that require an Expert User for identification are queued up for the Expert User(s).

[0103] FIG. 5 illustrates one embodiment of New Object Analysis 460. As shown, first it is determined at step 510 whether or not there were any matches that exceeded a minimum "close match" confidence threshold. If no matches exceed the close match confidence threshold, then the system assumes that the unidentified object in the image is an object currently unknown to the system.

[0104] The close match threshold is typically lower than the best match confidence threshold used in the Best Match Analysis, discussed above, although in one embodiment, the same threshold configuration may be used for each process. Alternatively, a close match minimum threshold may not be set at all, and the process always continues directly to step 520.

[0105] In the embodiment shown in FIG. 5, if it is determined at step 510 that there was at least one profile whose confidence score exceeded the close match minimum confidence threshold, the system assumes that the unidentified object in the image set is an object associated with one of the matches that exceeds the threshold. In this case, the process continues to step 520. Because multiple profiles can be stored for a single object, and therefore a ranked list of matches returned at step 440 can include multiple profiles that are associated with the same object, at step 520, the list of matches is examined to eliminate duplicate objects in the list. That is, a list of candidate objects is determined.

[0106] For example, suppose three confidence scores exceed the close match minimum threshold of 85 - Art with a first score of 88 corresponding to a first profile associated with Art, and a second score of 91 corresponding to a second profile associated with Art, and Bert with a score of 93 for the only profile associated with Bert - then the list of candidate objects includes Art and Bert, and one of the instances of Art may be removed from the list.

[0107] At step 525 each candidate object determined in step 520 is presented as a possible candidate to the Expert User. In one embodiment, the surveillance system image identified at step 405 is shown to the Expert User along with one of the stored images associated with the candidate. In the example given above, the Expert User may be presented with the identified surveillance system image and a single image of Art and a

single image of Bert. Alternatively, the Expert User may be presented with a video clip from which the images were obtained, or other associated video clips.

[0108] The Expert User determines whether the object in the identified image is one of the candidates. If the object is one of the candidates, the Expert User confirms this as a positive identification. The system will then either save the profile created at step 410, adding it as a new profile associated with the person, or discard the profile.

[0109] The decision of whether to save or discard the new profile can be made in a number of ways. The Expert User may be asked to manually confirm that the profile should be saved. Alternatively, the system may calculate the number of profiles that have already been saved for an object and discard the new profile if the number of profiles exceeds a certain number. Alternatively, the new profile may be saved, while an older profile is discarded, in order to maintain a consistent number of profiles for a person. In yet another alternative, the confidence score may be used as a determining factor in whether or not to save the profile. Many alternatives will be apparent to those skilled in the art.

[0110] In this process, the Expert User views the image of the object as captured in the surveillance video, and provides a name of the person and/or other information that identifies the object to the system. The system stores the object identifying information in a database. In addition, the profile created in step 410 and, optionally, an image of the object acquired by the surveillance system (such as an image from the image set identified in step 405) are saved in appropriate databases and associated with the object identifying information. Any method that saves the profile extracted in step 410 and associates it with a person or object can be used.

[0111] It is possible that the Expert User will not be able to identify the person or object in the video surveillance image identified in step 405. In this case, the profile could be discarded. Alternatively, the video clip and/or image acquired by the surveillance system and profile extracted for that image could be stored as an "unidentified" person or object. In this case, the person or object could be identified at a later time when more information is available, or flagged as a person of interest. In addition, the Expert User may be allowed to match the profile with another object in the database that was not in the list of candidate objects.

[0112] Multiple profiles associated with a single object improve the recognition accuracy rate. However, it may be impossible or impractical to save every profile associated with an object. Therefore, embodiments of the present invention will discard certain profiles in order to minimize the number of profiles saved per object. Profiles are typically saved only if they add value. For example, if a new profile has nearly identical measurements to an existing profile, then one of the two profiles may be discarded. However, if a new profile has significantly different measurements than the measurements in all previously existing profiles of the same object, then the new profile may be retained.

[0113] Furthermore, in alternative embodiments, the Expert User does not have to be a human. For instance, objects can be recognized using "expert" identification techniques. Such expert identification techniques may be too computationally expensive to be practical for the initial identification operation. While it may be impractical to invoke such techniques every time an identification operation is performed, it may be practical to invoke such techniques for the relative fewer situations in which the initial identification operation fails to identify the object of an image. In one embodiment, feedback from the expert user's identifications is used to "train" the system.

[0114] In addition, prior Expert User identifications or corrections to identifications made by the system can be used as a factor in the best match analysis to improve recognition. For example, prior Expert User corrections can be used to weight confidence scores, or as a factor when determining a weighted average in Best Match Analysis.

[0115] For example, a running score of correct and incorrect matches made from each stored profile may be kept. Any future matches using that profile may then be weighted accordingly. Furthermore, if the percentage of incorrect matches associated with a particular profile is determined to be too high, the profile may be "pruned" from the set of stored profiles. Likewise, a profile that is associated with a high correct identification rate may be weighted higher in the Best Match Analysis.

[0116] Embodiments of the present invention allow for reliable face recognition technology in an everyday type of environment, such as a surveillance system in an office building. Over time, multiple profiles are confirmed for a single person, thereby providing for higher accuracy rates as the system is used. The more samples (i.e., profiles) there are associated with a person, the more likely that the system will correctly identify people. By

going through the above described process, the system is "trained" to more accurately identify people.

GENERAL COMPUTER SYSTEM

[0117] FIG. 2 is a block diagram that illustrates a computer system 200 upon which an embodiment of the invention may be implemented. Computer system 200 includes a bus 202 or other communication mechanism for communicating information, and a processor 204 coupled with bus 202 for processing information. Computer system 200 also includes a main memory 206, such as a random access memory (RAM) or other dynamic storage device, coupled to bus 202 for storing information and instructions to be executed by processor 204. Main memory 206 also may be used for storing temporary variables or other intermediate information during execution of instructions to be executed by processor 204. Computer system 200 further includes a read only memory (ROM) 208 or other static storage device coupled to bus 202 for storing static information and instructions for processor 204. A storage device 210, such as a magnetic disk or optical disk, is provided and coupled to bus 202 for storing information and instructions.

[0118] Computer system 200 may be coupled via bus 202 to a display 212, such as a cathode ray tube (CRT), for displaying information to a computer user. An input device 214, including alphanumeric and other keys, is coupled to bus 202 for communicating information and command selections to processor 204. Another type of user input device is cursor control 216, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to processor 204 and for controlling cursor movement on display 212. This input device typically has two degrees of freedom in two axes, a first axis (e.g., x) and a second axis (e.g., y), that allows the device to specify positions in a plane.

[0119] The invention is related to the use of computer system 200 for implementing the techniques described herein. According to one embodiment of the invention, those techniques are performed by computer system 200 in response to processor 204 executing one or more sequences of one or more instructions contained in main memory 206. Such instructions may be read into main memory 206 from another machine-readable medium, such as storage device 210. Execution of the sequences of instructions contained in main memory 206 causes processor 204 to perform the process steps described herein. In

alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions to implement the invention. Thus, embodiments of the invention are not limited to any specific combination of hardware circuitry and software.

[0120] The term "machine-readable medium" as used herein refers to any medium that participates in providing data that causes a machine to operation in a specific fashion. In an embodiment implemented using computer system 200, various machine-readable media are involved, for example, in providing instructions to processor 204 for execution. Such a medium may take many forms, including but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media includes, for example, optical or magnetic disks, such as storage device 210. Volatile media includes dynamic memory, such as main memory 206. Transmission media includes coaxial cables, copper wire and fiber optics, including the wires that comprise bus 202. Transmission media can also take the form of acoustic or light waves, such as those generated during radio-wave and infra-red data communications.

[0121] Common forms of machine-readable media include, for example, a floppy disk, a flexible disk, hard disk, magnetic tape, or any other magnetic medium, a CD-ROM, any other optical medium, punchcards, papertape, any other physical medium with patterns of holes, a RAM, a PROM, and EPROM, a FLASH-EPROM, any other memory chip or cartridge, a carrier wave as described hereinafter, or any other medium from which a computer can read.

[0122] Various forms of machine-readable media may be involved in carrying one or more sequences of one or more instructions to processor 204 for execution. For example, the instructions may initially be carried on a magnetic disk of a remote computer. The remote computer can load the instructions into its dynamic memory and send the instructions over a telephone line using a modem. A modem local to computer system 200 can receive the data on the telephone line and use an infra-red transmitter to convert the data to an infra-red signal. An infra-red detector can receive the data carried in the infra-red signal and appropriate circuitry can place the data on bus 202. Bus 202 carries the data to main memory 206, from which processor 204 retrieves and executes the instructions. The instructions received by main memory 206 may optionally be stored on storage device 210 either before or after execution by processor 204.

[0123] Computer system 200 also includes a communication interface 218 coupled to bus 202. Communication interface 218 provides a two-way data communication coupling to a network link 220 that is connected to a local network 222. For example, communication interface 218 may be an integrated services digital network (ISDN) card or a modem to provide a data communication connection to a corresponding type of telephone line. As another example, communication interface 218 may be a local area network (LAN) card to provide a data communication connection to a compatible LAN. Wireless links may also be implemented. In any such implementation, communication interface 218 sends and receives electrical, electromagnetic or optical signals that carry digital data streams representing various types of information.

[0124] Network link 220 typically provides data communication through one or more networks to other data devices. For example, network link 220 may provide a connection through local network 222 to a host computer 224 or to data equipment operated by an Internet Service Provider (ISP) 226. ISP 226 in turn provides data communication services through the world wide packet data communication network now commonly referred to as the "Internet" 228. Local network 222 and Internet 228 both use electrical, electromagnetic or optical signals that carry digital data streams. The signals through the various networks and the signals on network link 220 and through communication interface 218, which carry the digital data to and from computer system 200, are exemplary forms of carrier waves transporting the information.

[0125] Computer system 200 can send messages and receive data, including program code, through the network(s), network link 220 and communication interface 218. In the Internet example, a server 230 might transmit a requested code for an application program through Internet 228, ISP 226, local network 222 and communication interface 218.

[0126] The received code may be executed by processor 204 as it is received, and/or stored in storage device 210, or other non-volatile storage for later execution. In this manner, computer system 200 may obtain application code in the form of a carrier wave.

[0127] In the foregoing specification, embodiments of the invention have been described with reference to numerous specific details that may vary from implementation to implementation. Thus, the sole and exclusive indicator of what is the invention, and is intended by the applicants to be the invention, is the set of claims that issue from this

application, in the specific form in which such claims issue, including any subsequent correction. Any definitions expressly set forth herein for terms contained in such claims shall govern the meaning of such terms as used in the claims. Hence, no limitation, element, property, feature, advantage or attribute that is not expressly recited in a claim should limit the scope of such claim in any way. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

CLAIMS

What is claimed is:

- 1 1. A method of determining a best match between a target profile and a set of stored
2 profiles, wherein a profile contains a set of identifying information extracted from an image
3 set associated with the profile, the method comprising computer-implemented steps of:
4 generating a plurality of confidence scores based on comparisons between the target
5 profile and the set of stored profiles;
6 weighting the generated confidence scores using information external to the
7 confidence scores; and
8 selecting a stored profile as the best match for the target profile based on the plurality
9 of weighted confidence scores.
- 1 2. The method of Claim 1 wherein each stored profile has an age, and the information
2 used to weight the confidence scores comprises profile age information.
- 1 3. The method of Claim 2 wherein a first stored profile is older than a second stored
2 profile, and wherein a weighting factor for said second stored profile is greater than a
3 weighting factor for said first stored profile.
- 1 4. The method of Claim 1 wherein the image set associated with the target profile is
2 captured by a first camera at a first time, and wherein the information used to weight the
3 confidence scores comprises information associated with whether a best matching profile has
4 been determined for another profile associated with another image set captured by the first
5 camera within a predetermined time frame of the first time.
- 1 5. The method of Claim 1 wherein the image set associated with the target profile is
2 captured by a first camera at a first time, and the information used to weight the confidence
3 scores comprises information associated with whether a best matching profile has been
4 determined for another profile associated with another image set captured by a second
5 camera within a predetermined time frame of the first time.

1 6. The method of Claim 1 wherein the information used to weight the confidence scores
2 comprises a rate of confirmed correct identifications associated with at least one profile in the
3 set of stored profiles.

1 7. The method of Claim 6 wherein a first stored profile has a confirmed correct
2 identification rate that is greater than a confirmed correct identification rate of a second
3 stored profile, and wherein a weighting factor for said first stored profile is greater than a
4 weighting factor for said second stored profile.

1 8. The method of Claim 1 wherein an image set is acquired from a multi-camera video
2 surveillance system.

1 9. A method for maintaining associations between profiles and objects in an object
2 recognition system, wherein a profile comprises a set of identifying information extracted
3 from an image set associated with the profile, comprising the computer-implemented steps
4 of:

5 automatically creating a first association between a first stored profile and a first
6 object;
7 automatically creating a second association between a second stored profile and the
8 first object;
9 providing a view of the image sets associated with the first and second stored profiles
10 to a user;
11 receiving feedback from the user about the second association; and
12 modifying the second association in accordance with the received feedback.

1 10. The method of Claim 9, wherein an object comprises a person, and a set of
2 identifying information comprises a facial feature set.

1 11. The method of Claim 9, wherein the step of providing a view of the image sets
2 associated with the first and second stored profiles additionally comprises providing a view
3 of the first and second associations to the user.

1 12. The method of Claim 11, wherein the step of receiving feedback includes receiving
2 feedback that the second association is correct, and wherein the step of modifying includes
3 storing confirmation information with the second association.

1 13. The method of Claim 11, wherein the step of receiving feedback includes receiving
2 feedback that the second association is incorrect, and the step of modifying includes
3 removing the second stored profile from the object recognition system.

1 14. The method of Claim 11, wherein the step of receiving feedback includes receiving
2 feedback indicating that the second stored profile should be associated with a second object;
3 and the step of modifying includes creating an association between the second stored profile
4 and the second object.

1 15. The method of Claim 9, further comprising the computer-implemented steps of:
2 receiving a new image set;
3 extracting a set of identifying information from the new image set;
4 creating a new profile that includes the extracted set of identifying information;
5 providing a view of the new image set to a user;
6 receiving feedback from the user indicating an object that should be associated with
7 the new profile; and
8 creating a new association between the new profile and the object indicated by the
9 user.

1 16. The method of Claim 15, wherein the step of receiving feedback from the user
2 indicating an object comprises receiving feedback from the user indicating that the new
3 profile should be associated with the first object; and wherein the step of creating an
4 association comprises creating an association between the new profile and the first object,
5 such that the first stored profile, the second stored profile and the new profile are all
6 associated with the first object.

1 17. The method of Claim 15, wherein the step of providing a view of the new image set
2 additionally comprises providing a list of potential object identifications for the new profile,

3 and the step of receiving feedback from the user indicating an object comprises receiving a
4 choice of one object from the list.

1 18. The method of Claim 15, wherein the step of receiving feedback from the user
2 indicating an object comprises receiving feedback that a new object should be created, and
3 the step of creating a new association comprises creating a new object and creating an
4 association between the new profile and the new object.

1 19. A method of determining a best matching object for a target profile in an object
2 recognition system in which each object recognized by the system is associated with a
3 plurality of stored profiles, wherein a profile contains a set of identifying information
4 extracted from an image set associated with the profile, the method comprising the computer-
5 implemented steps of:

6 generating a plurality of confidence scores based on comparisons between the target
7 profile and a set of stored profiles;
8 determining an object associated with each generated confidence score;
9 analyzing the generated confidence scores and determined associated objects; and
10 selecting a best matching object in accordance with the analyzed confidence scores
11 and determined associated objects.

1 20. The method of Claim 19, wherein the step of analyzing comprises determining, for
2 each object, a number of occurrences of that object being associated with a confidence score,
3 and the step of selecting a best matching object comprises selecting the object with a greatest
4 number of occurrences.

1 21. The method of Claim 20, wherein the step of analyzing comprises determining, for
2 each object, a number of occurrences of that object being associated with a confidence score
3 that exceeds a predetermined minimum threshold.

1 22. The method of Claim 19, wherein the step of analyzing comprises calculating, for
2 each object, a weighted average of confidence scores associated with that object; and the step
3 of selecting a best matching object comprises selecting the object with a greatest weighted
4 average.

- 1 23. The method of Claim 22, wherein each stored profile has an age, wherein the step of
2 calculating a weighted average of confidence scores comprises:
3 for each object, weighting each confidence score associated with the object in
4 accordance with the age of the profile associated with the confidence score, and
5 averaging all weighted confidence scores associated with that object.
- 1 24. The method of Claim 22, wherein the step of calculating a weighted average of
2 confidence scores comprises:
3 for each object, determining when the object was last identified by the system, and
4 weighting the average confidence scores associated with that object in accordance
5 with when the object was last identified.
- 1 25. The method of claim 22, additionally comprising the step of:
2 ranking the generated confidence scores;
3 wherein the step of calculating a weighted average of confidence scores comprises
4 using the rankings to weight confidence scores.
- 1 26. The method of claim 22, additionally comprising the step of:
2 receiving access information from an access control system;
3 wherein the step of calculating a weighted average of confidence scores comprises
4 using the received access information to weight the confidence scores.
- 1 27. The method of claim 22, wherein an object comprises a person, and a set of
2 identifying information comprises a facial feature set.
- 1 28. The method of claim 27, additionally comprising the step of:
2 for each person associated with a profile used to generate a confidence score,
3 determining organizational information about the person,
4 wherein the step of calculating a weighted average of confidence scores comprises
5 using the organizational information to weight the confidence scores.
- 1 29. The method of claim 28, wherein the target profile has a creation time, and the
2 organizational information includes a work shift associated with the person, wherein the step

3 of calculating a weighted average of confidence scores comprises determining whether the
4 work shift of the person corresponds to the creation time of the target profile, and using said
5 determination to weight the confidence scores.

1 30. The method of claim 28, wherein the target profile has a location, and the
2 organizational information includes at least one authorized location associated with the
3 person, wherein the step of calculating a weighted average of confidence scores comprises
4 determining whether any authorized location of the person corresponds to the location of the
5 target profile, and using said determination to weight the confidence scores.

1 31. The method of claim 28, wherein the target profile has a location, and the
2 organizational information includes at least one organizational department associated with
3 the person, wherein the step of calculating a weighted average of confidence scores
4 comprises determining whether any organizational department of the person corresponds to
5 the location of the target profile, and using said determination to weight the confidence
6 scores.

1 32. The method of claim 28, wherein the organizational information includes at least one
2 organizational department associated with the person, wherein the step of calculating a
3 weighted average of confidence scores comprises determining whether another person
4 associated with at least one organization department associated with the person has been
5 identified by the system in time and location proximity with the person.

1 33. The method of Claim 22, wherein each stored profile has a rate of correct
2 identifications, wherein the step of calculating a weighted average of confidence scores
3 comprises:
4 for each object, weighting each confidence score associated with the object in
5 accordance with the correct identification rate of the profile associated with the confidence
6 score, and averaging all weighted confidence scores associated with that object.

1 34. The method of Claim 19, additionally comprising the steps of:

2 determining if any of the plurality of confidence scores exceed a predetermined
3 minimum threshold, and if no scores exceed the minimum threshold, then the method further
4 comprising the steps of:
5 providing a view of the image set associated with the target profile to a user to
6 determine if the object in the image set matches an object recognized by the object
7 recognition system, and if the user determines that the object in the image does not match a
8 recognized object, enrolling the object as a new object and saving the target profile
9 associated with the new object.

1 35. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to
3 perform the method recited in Claim 1.

1 36. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to
3 perform the method recited in Claim 9.

1 37. A computer-readable medium carrying one or more sequences of instructions which,
2 when executed by one or more processors, causes the one or more processors to
3 perform the method recited in Claim 19.




FIG. 1





FIG. 2

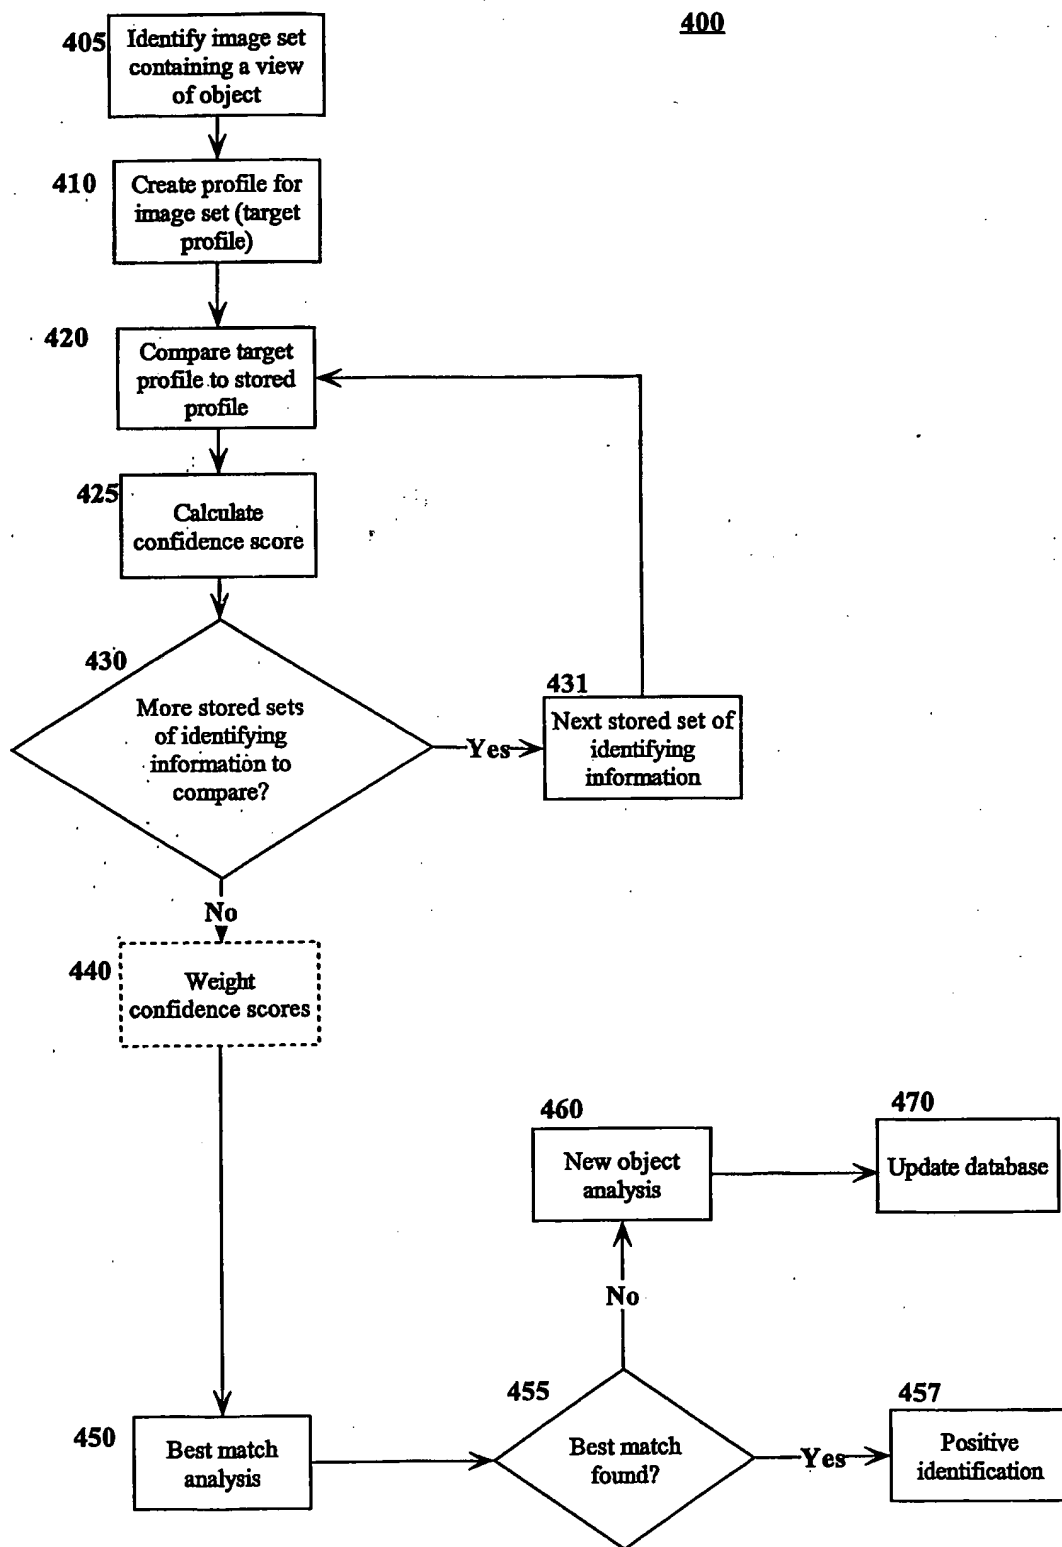


300

FIG. 3

4/6

FIG. 4



5/6
FIG. 5

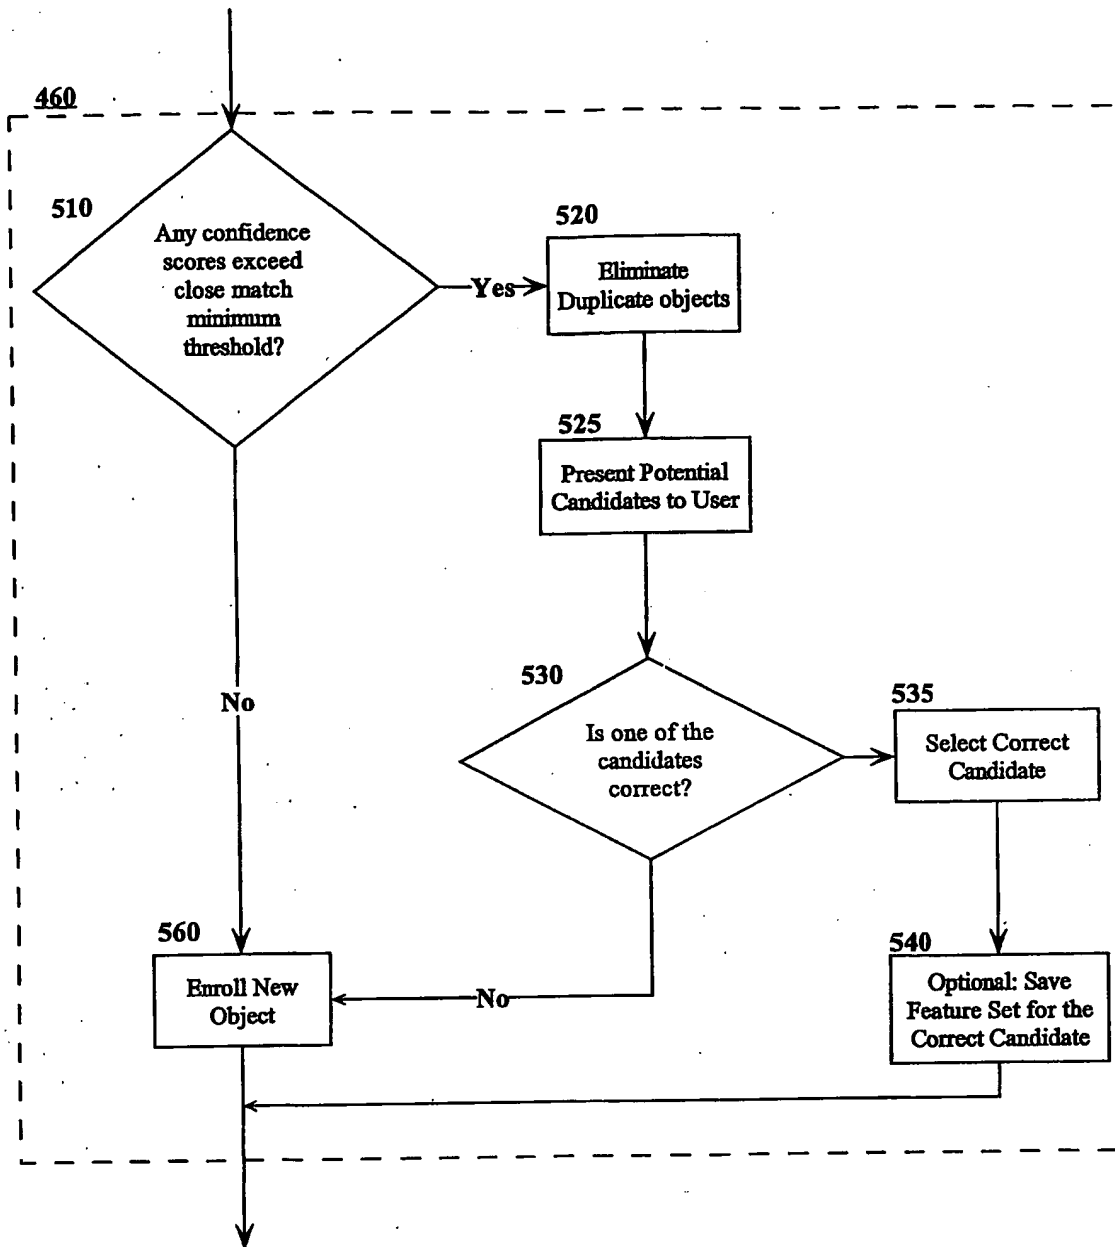



FIG. 6



INTERNATIONAL SEARCH REPORT

Application No
PCT/US2005/008735

A. CLASSIFICATION OF SUBJECT MATTER IPC 7 G06K9/00				
According to International Patent Classification (IPC) or to both national classification and IPC				
B. FIELDS SEARCHED Minimum documentation searched (classification system followed by classification symbols) IPC 7 G06K				
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched				
Electronic data base consulted during the international search (name of data base and, where practical, search terms used) EPO-Internal, INSPEC, COMPENDEX				
C. DOCUMENTS CONSIDERED TO BE RELEVANT				
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No		
X	OKADA K ET AL: "Automatic video indexing with incremental gallery creation: integration of recognition and knowledge acquisition" KNOWLEDGE-BASED INTELLIGENT INFORMATION ENGINEERING SYSTEMS, 1999. THIRD INTERNATIONAL CONFERENCE ADELAIDE, SA, AUSTRALIA 31 AUG.-1 SEPT. 1999, PISCATAWAY, NJ, USA, IEEE, US, 31 August 1999 (1999-08-31), pages 431-434, XP010370975 ISBN: 0-7803-5578-4 page 433, left-hand column, last paragraph --- -/--	9-11, 15-18		
<input checked="" type="checkbox"/> Further documents are listed in the continuation of box C. <input type="checkbox"/> Patent family members are listed in annex				
* Special categories of cited documents : <table border="0"> <tr> <td> "A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier document but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed </td> <td> "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance, the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance, the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art "B" document member of the same patent family </td> </tr> </table>			"A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier document but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance, the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance, the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art "B" document member of the same patent family
"A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier document but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance, the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance, the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art "B" document member of the same patent family			
Date of the actual completion of the international search 13 July 2005		Date of mailing of the international search report 05/09/2005		
Name and mailing address of the ISA European Patent Office, P B 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel (+31-70) 340-2040, Tx 31 651 epo nl, Fax (+31-70) 340-3016		Authorized officer Granger, B		

INTERNATIONAL SEARCH REPORT

Application No
PCT/US2005/008735

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No
X	<p>S. GONG ET AL: "Dynamic Vision : from Images to Face Recognition" 2000, IMPERIAL PRESS , LONDON, GB , XP002335641 page 16, last paragraph page 276 page 277, line 25 - line 27 page 286, line 21 - line 22 page 288, line 17 - line 18 page 288, line 22 - line 25 page 291, line 25 - line 26 page 293, line 10 - line 12 page 293, line 31 - line 32</p>	<p>9-11, 15-18</p>
X	<p>RUUD BOLLE ET AL: "Guide to Biometrics" 2003, SPRINGER , NEW YORK, USA , XP002335642 page 145, paragraph 5 - paragraph 7 page 159 - page 161 page 169 - page 170</p>	<p>9-11, 15-18</p>

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2005/008735

Box II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This International Search Report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.: 1-8, 12-14, 19-37
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☒ Claims Nos. 1-8, 12, 14, 19-37
because they relate to parts of the International Application that do not comply with the prescribed requirements to such an extent that no meaningful International Search can be carried out, specifically:
see FURTHER INFORMATION sheet PCT/ISA/210
3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

1. ☐ As all required additional search fees were timely paid by the applicant, this International Search Report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fee, this Authority did not invite payment of any additional fee.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this International Search Report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this International Search Report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest.
☐ No protest accompanied the payment of additional search fees

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

Continuation of Box II.2

Claims Nos.: 1-8,12,14,19-37

1 The subject-matter of claims 1-8, 19-35, 37 is not disclosed in a manner sufficiently clear and complete for the claimed invention to be carried out by a person skilled in the art, in violation of Article 5 PCT. Hence the International Preliminary Examining Authority considers that the international application fails to comply with the prescribed requirements of Article 5 PCT to such an extent that no meaningful opinion can be formed on the novelty, or inventive step of the claimed invention under Article 17(2)(a)(ii) PCT.

1.1 As the subject-matter of claims 1-8, 19-35, 37 deals with computations performed on "confidence scores based on comparaisons between the target profile and the set of stored profiles", determining these confidence scores is essential for the claimed invention.

1.2 However the application does not disclose how a confidence score is computed for each comparaison.

1.3 Furthermore the basic handbooks in the field of pattern recognition do not disclose how to compute a confidence score for a comparaison. Hence determining confidence scores for comparaisons is not included in the common knowledge of the person skilled in the art.

1.4 Consequently the application does not contain sufficient information to allow the person skilled in the art, using his common knowledge, to perform the claimed invention without undue burden.

2 The subject-matter of claims 12-14 is not disclosed in a manner sufficiently clear and complete for the claimed invention to be carried out by a person skilled in the art, in violation of Article 5 PCT. Hence the International Preliminary Examining Authority considers that the international application fails to comply with the prescribed requirements of Article 5 PCT to such an extent that no meaningful opinion can be formed on the novelty, or inventive step of the claimed invention under Article 17(2)(a)(ii) PCT.

2.1 These claims relate to the modification of an "association" between a "profile" and an "object", which association has already been created (and thus stored) in the computer (see claim 1). In claim 12 (resp. 13; 14) it is not clear how the association is modified by storing confirmation information (resp. by removing the profile — how can the association still exist?; by creating an association which supposedly already exists).

3 The International Preliminary Examining Authority considers that the subject-matter of claim 36 relates to a subject-matter on which the International Preliminary Examining Authority is not required, under the Regulations, to carry out an international preliminary examination, and

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

in this particular case decides not to carry out such examination under Article 17(2)(a)(i).

3.1 The subject-matter of claim 36 does not relate to an apparatus, a process to control the apparatus, a product made by the apparatus or a use of this product, but instead to a "medium carrying one or more sequences of instructions", Rule 39.1.

3.2 This claim is understood as being related to a computer program (the latter being defined as a "sequence of computer instruction". As the International Preliminary Examining Authority does not have access to any complete and searchable collection of computer programs (programs either in the form of computer program code or in the form of executable computer software), the Authority is not equipped to search prior art concerning such programs under Rule 39(1)(vi) PCT.

The applicant's attention is drawn to the fact that claims relating to inventions in respect of which no international search report has been established need not be the subject of an international preliminary examination (Rule 66.1(e) PCT). The applicant is advised that the EPO policy when acting as an International Preliminary Examining Authority is normally not to carry out a preliminary examination on matter which has not been searched. This is the case irrespective of whether or not the claims are amended following receipt of the search report or during any Chapter II procedure. If the application proceeds into the regional phase before the EPO, the applicant is reminded that a search may be carried out during examination before the EPO (see EPO Guideline C-VI, 8.5), should the problems which led to the Article 17(2) declaration be overcome.